



## EMC's Isilon and Its Role(s) for Big Data

Analyst: Jim Baker

### Management Summary

In December 2010, EMC completed the purchase of Isilon Systems, Inc., of Seattle and since has announced various product line extensions and multiple user success stories. EMC's Isilon is a leading provider of scale-out NAS solutions across many industries. Among Isilon's key claims to fame are its *OneFS* file system, commodity-based appliances complete with storage, and auxiliary software packages that are especially appropriate for working with so-called "Big Data." Was this a defensive acquisition so that some other vendor would not acquire Isilon for their portfolio or was it more strategic in nature – filling out a hole in the EMC range of capabilities in an interesting future-oriented business segment? It is the latter and it should be increasingly obvious that the future is fast approaching (or has arrived) when Big Data problems and solution sets should be on most everyone's radar, including yours.

So, what is Big Data anyway? In some ways, the "big" in Big Data is in the eyes of the beholder. Most experts agree that the threshold for Big Data is not reached until there are multi-terabytes or exabytes or soon will be zettabytes to store and analyze. In practical terms, data becomes "big" when datasets are so large that they become difficult to manipulate in real time with standard IT tools, such as file systems, relational database software, analytical tools, and common server and storage hardware and software platforms. "Big" really means "Too Big"!

Certain applications and industries organically make use of Big Data. Among these are oil and gas exploration, weather forecasting and climate change analysis, healthcare (medical records, and image-based files, such as MRIs and Cat Scans), genomics, digital media and entertainment, telephony, online retailing, Internet publishing, and more. Other Big Data applications come into existence not from the nature of the data being kept, but rather from the sheer volume of the data and the length of time the data needs to be available to the user community.

Another use case for Big Data involves how quickly the data must be processed and acted upon with the compounding factor being how much data is being ingested at any point in time. An example would be the processing of surveillance data from multiple points simultaneously, where an active (real-time) response often is required. Think of feeds from multiple sensors surrounding a nuclear power plant. Intrusion detection requires a multipronged response yet most of the time the system processes a large amount of innocuous data. A second example is fraud detection during periods of high transaction volumes. An organized attack on a financial institution must have immediate and multiple reactions. In both examples, processing occurs immediately upon capture. To learn more about how EMC's Isilon product family is able to data capture, store, process, and analyze Big Data, please read on.

### IN THIS ISSUE

➤ <b>Virtuous/Vicious Cycles</b> .....	<b>2</b>
➤ <b>Isilon's Big Data Virtuous Cycle</b> .....	<b>2</b>
➤ <b>The Isilon Product Family</b> .....	<b>2</b>
➤ <b>Ten Reasons Why You Should Consider Isilon</b> .....	<b>5</b>
➤ <b>Conclusion</b> .....	<b>6</b>

## Virtuous and Vicious Cycles

A *virtuous cycle* is an economic concept whereby one action positively reinforces itself by means of a continuous feedback loop. Each iteration or result is ever-increasing or positive unless or until an event or a new technology breaks the cycle disruptively. Conversely, a *vicious cycle* is one that negatively reinforces itself by means of a feedback loop and is ever-decreasing or continually detrimental until its cycle is broken. Because they can be a source of long-term revenues and high consumer loyalty, savvy businesses always are on the lookout for a virtuous cycle that involves the increased use of its technology either as a service or as products that support the cycle.

An example from the package distribution industry is instructive. Consider how the presence of FedEx Corporation forever changed the package delivery business. Beginning with simple overnight document delivery, FedEx used a novel (at the time) hub and spoke aircraft routing system. All packages flew inbound to Memphis, Tenn., where they are sorted and put on outbound planes heading to distant destinations. This obviated the necessity for multiple point-to-point airline routes and was highly efficient. Over time, FedEx grew their business to move packages as well and to offer different delivery options beyond just overnight delivery. Sharp businesses saw the FedEx pipeline as a way to implement *Just-In-Time* supply chains, thereby reducing the need for stockpiles of products or spare parts in many expensive leased buildings around the world. Less brick. Less mortar. Less hassle. Yet repair parts are available (the next day), when and where they are needed. As the new “supply chain experience” worked efficiently with few or no hiccups, businesses now can commit to similar strategies for additional product lines thereby further increasing the virtuous cycle. EMC’s Isilon represents a solution with just such a virtuous cycle opportunity.

## Isilon’s Big Data Virtuous Cycle

The Isilon logo of two apparently moving ovals where one seems to be sustaining the other and where both are “in motion” is an insightful graphical representation for Isilon’s Big Data Virtuous Cycle.

- Capture original data and store it for future analysis.
- Analyze it to gain insights.
- Capture and store incremental data.
- Aggregate and re-



analyze it to gain additional insights (including comparisons before and after incremental data).

- Add sensors to gather more data, if appropriate and/or add additional storage, as required.
- Re-analyze data to gain *even more* insights leading to increased value of data to the organization.
- Repeat.

Each builds on the other, learning occurs throughout the process, and feedback is implicit in the constant march forward – all classical characteristics of an economic virtuous cycle.

## The Isilon Product Family

### OneFS Operating System

OneFS, a modern operating system for file-based (and not block) data, does not depend upon specialized hardware as a critical part of the storage architecture. It combines a file system, volume manager, and RAID protection into a single entity that spans all of the physical nodes within the storage platform. Its features include data protection and automated data balancing and migration, along with the ability to add storage and performance capabilities without taking down system applications. This latter point is hugely important in the world of ever-growing data.

OneFS works with one or more Isilon clusters. A cluster has multiple storage “nodes”, which are rack-mounted modules containing memory, CPU, networking, nonvolatile RAM, Infiniband interconnects, as well as storage media itself. Cache, capacity, and performance from all nodes are shared across the whole system under one level of management. A typical cluster contains three to 144 nodes. As additional nodes are added, all resources of each node are accumulated and made immediately available for the benefit of the system as a whole. Data movement within the storage system is done across the Infiniband interconnects, thereby avoiding important user and application network traffic. Performance is enhanced by the parallelism implicit in Isilon’s multi-node architecture. Data protection is included in the system using a feature called *FlexProtect*. FlexProtect, based on Reed-Solomon encoding technology, provides protection for up to four simultaneous failures of either full nodes or individual drives. Reconstruction of failed modules also takes advantage of the system’s massive parallelism. Hot spares are virtual in that they can be anywhere in the system where space is available. Identification and isolation of failed modules is constantly

ongoing, with transparent relocation of data from failed modules to good ones.

### **Hardware Platforms**

#### **S-Series Platform Node**

This purpose built hardware platform is for high I/O (Input/Output) intensive applications, including transaction processing. Using Isilon's scale-out design based on parallelism, this module enables high transaction rates for file-based applications currently residing on high-performance storage. The *S-Series* node supports SSDs for metadata operations and other resource-intensive tasks. Like other Isilon hardware products, the S-Series is designed for on the fly installation of additional nodes with no application interruption. A wide range of storage options allows customer choice – in a mix and match fashion. Backend data transfers are over the private, high-performance InfiniBand network.

#### **X-Series Platform Node**

The *X-Series* strikes a balance between large capacity and high performance storage. It too supports SSD technology and has a wide range of storage options. The InfiniBand interconnect is standard. Use the X-Series for virtualized environments, high-performance computing support, and most file based applications. Tier-2 or Tier-3 data is typically stored on the X-Series because it is not often being used transactionally.

#### **NL-Series Platform Node**

The NL-Series provides lower cost, yet still scalable nearline storage under a single global namespace. Its use cases are most appropriate for disk-based online archiving and disaster recovery, where most of the data is not likely to be retrieved (very often). The NL-Series has the lowest performance profile of all three and, as such, has the lowest price.

#### **Pricing Model**

For an actual quotation for your specific needs, you must contact EMC. However, this is how pricing is established. For the S-Series and X-Series, which allow users to choose custom drive configurations, nodes are priced based on the drive configuration (e.g., a node with no SSDs will cost less than one with six). On a system basis, nodes are sold on a per node pricing basis. You must buy three nodes with your first purchase to create a minimal cluster but, thereafter, nodes can be purchased one at a time or many at a time, and are priced per node.

### **Hardware Extension Nodes**

#### **Capacity Extension Nodes**

These nodes simply add more storage space to

the Isilon system without downtime and vary according to the disk technology being added. These nodes use standard NAS file sharing protocols over Gigabit Ethernet and offer full support for *SyncIQ* asynchronous replication software. More on SyncIQ follows.

#### **Performance Extension Nodes**

*IQ Accelerators* are performance-enhancing nodes that can be added to either the S-Series or the X-Series cluster when additional throughput is required. IQ Accelerators increase performance independent of type or speed of storage in storage nodes. Multi-Protocol Support includes NFS v3 (UDP or TCP), CIFS, HTTP, FTP, NDMP, SNMP, LDAP, ADS, and NIS reads/writes. (For a full list, contact EMC.)

*Backup Accelerators* add up to four concurrent streams to the cluster and have been certified for use with most industry standard backup software packages from Symantec, IBM, CommVault, EMC, Atempo, and BakBone. (For a full list, contact EMC.)

### **Enterprise Application Software**

#### **SmartPools**

*SmartPools* allow a single file system for multiple tiers with automated real-time data movement controlled by user policy. SmartPools are for use with S-Series, X-Series, and NL-Series. With SmartPools, a single file system easily is deployed to span multiple tiers of performance and capacity. Workflows seamlessly adapt to process changes without affecting applications or other workflows. No application changes are necessary. SmartPools can use solid-state drives (SSDs) as a tier for metadata, spread across all drives in the system to enable exceptionally high performance.

#### **SmartQuotas**

*SmartQuotas* seamlessly partition storage into easily managed segments for quotas at the cluster, directory, sub-directory, user, and group levels. SmartQuotas is the name for the feature that enables thin provisioning across the system. Thin provisioning lets applications share storage space in a manner that the actual consumption of space does not occur until something is actually stored. Meanwhile, other applications can share pool space and grow as they, too, need it. Thin provisioning has been widely credited with improving the efficiency of contemporary storage (by eliminating the excess storage allocated to provide enough "headroom" to be safe). Unique quota and provisioning policies can be set at any level. Reporting and notification functions can be tailored to site-specific requirements. Users may select

from *hard* (immediate application interruption), *soft* (pending application interruption if no actions taken), and/or *advisory* (general warning or rule) limits set across the organization for specific users and groups, and across various directory structures.

#### SnapshotIQ

*SnapshotIQ* enables locally-retained, read-only snapshots as incorruptible point-in-time images of data. These snapshots are spread across the parallel structures of Isilon storage nodes. Snapshots are highly space efficient because they capture changes only and avoid the overhead of saving multiple full images. When used for data protection purposes, SnapshotIQ allows users to recover from accidental deletions, corruptions, or modifications to their data. Because they incur minimal performance overhead and can be created safely on in-production systems, users can make frequent user-recoverable images on demand.

#### Aspera for Isilon IQ

Aspera, a file transfer optimization company located in Emeryville, California, and Isilon Systems have jointly created a solution that enables easy distribution of large file sets called *Aspera Enterprise Server* for Isilon IQ. Functionally, this product collapses related server and storage tiers into a single, common storage, and transfer management pool whereby the combined pool can be moved in its entirety as a unit to another place or server. This product is aware of the huge files as well as the innate parallelism of Isilon. Performance and capacity can be added (or subtracted) without application downtime. No individual server or storage manipulations are required. Storage, bandwidth, and CPU utilization rates all are optimized because all nodes are working in parallel on individual or multiple simultaneous transfers.

#### InsightIQ

*InsightIQ* provides file system analytics and performance introspection within the Isilon family. Use cases include identification of performance bottlenecks in workflows and correlation and trend analysis on a per-directory or per-file basis. Use this feature as part of the Virtuous Cycle outlined above – to build out your actionable knowledge base. Users typically set up InsightIQ as a virtual appliance seamlessly integrated into enterprise-class environments.

#### SmartLock

*SmartLock* is a software based WORM (Write Once Read Many) solution. Mix and match SmartLock protected data with any other data on

the Isilon system. No extra hardware is required because this is a software-enabled feature. The unit of granularity for this protection is the file directory, which offers a huge advantage over implementations that provide WORM protection at the full volume level only.

#### SyncIQ

SyncIQ delivers extremely fast replication over the WAN or LAN. Because it uses the parallelism of the system, replication jobs are typically spread across the nodes of the system in an even distribution. SyncIQ is used for disaster recovery, disk-to-disk backup and restore, and parallelized content delivery tasks. Internal Isilon command and control functions utilize the InfiniBand interconnect, whereas transfers to systems outside the Isilon system use external network communications links.

#### SmartConnect

*SmartConnect* delivers intelligent client-connection load balancing of workloads and dynamic NFS failover and failback. Due to its architectural design, SmartConnect assures the user that even during a failure scenario, in-flight reads and writes will successfully completed. During a rare fail-over event, clients are distributed evenly across all surviving nodes in the cluster thereby providing minimal performance impact. The NFS failover support feature for Linux and UNIX clients is supported fully. SmartConnect is the technology that allows clients and storage nodes to be added (or subtracted) non-disruptively.

#### Cloud Neutral or Cloud Nothing

Isilon can be described as “cloud neutral”, meaning that it works equally well on either side of the corporate firewall and even could reside in a hybrid cloud, wherein elements of the solution exist simultaneously in both the private, as well as the public cloud.

Isilon also can be described as “cloud nothing” – when it is put into service in the classical IT way. In this case, classical means that the application is planned, procured, and managed by the internal IT department probably at the request of the business. It may include chargeback billings to departments consuming the IT service if that policy is implemented at the company. Regardless, all IT assets sit behind the corporate firewall, all data stored is proprietary to the corporation, and the IT infrastructure is not shared with any organization outside the company.

#### Public Cloud

Public cloud implementations implies that the infrastructure is being offered as a service by a

third party now referred to as an “Infrastructure as a Service” (IaaS) provider. Anyone or any firm in the public domain could procure these external services so there is implicit sharing of the infrastructure with all who “sign on.” This class of provider is a relatively recent phenomenon and is an outgrowth of “Software as a Service” provider offerings and other third-party offerings for specialized services, such as disaster recovery sites, backup and recovery sites, long-term storage sites, etc. Ownership of the assets remains with the provider who makes them available on a fee-for-service basis to clients of all sizes. Pricing is usually on a consumption basis, such that the client is paying only for what is being used and specifically not paying for standby or extra capacity. Isilon is a natural fit for public cloud implementations because of its scale-out infrastructure that expands storage space as more and more data is added to the repository.

#### Private Cloud

Private clouds are behind the corporate file wall yet they share several common characteristics with public clouds. They are elastic and can form, disintegrate, or re-form – as needed. They can be self-service once they have been designed and optimized for internal use by the corporate team. Their services are consumed as needed and they are priced (or charged back) to the ordering department. Maintenance often is performed by “lead operators”, who likely know the business cycle more than they know the underlying computing infrastructure. All of these features are designed in (or “architected”) by the corporate computing infrastructure staff, which ultimately are charged with arming users with “just right” technology to achieve business objectives easily and at the lowest reasonable cost. Again, Isilon is a natural fit here.

#### Hybrid Cloud

Hybrid clouds have elements of both private and public clouds. Applications can be spread across both domains mixing and matching the technologies to fit the task at hand. For long-term archiving, the company may procure space from a public cloud provider and place its data on shared storage likely after it has been encrypted for safety purposes. When and if it is needed, it can be recalled from the Public cloud and then unencrypted. Most recent data, on the other hand, might remain unencrypted behind the corporate firewall until it is no longer relevant or operationally necessary. Thereafter, the data could be demoted and placed in long-term archival storage out on the Isilon-based public cloud storage location.

## Ten Reasons Why You Should Consider Isilon

1. Isilon architecture and product features are “leading edge” in a world of rapidly-growing data.
2. The maximum size of their global namespace is greater than 15 PBs, which is likely more than sufficient for most applications. (Likewise, if it is not using all 15 PBs, then there is room for growth.) Isilon’s large namespace is the most meaningful characteristic of the product family, because it enables so many other uses. Clipper’s much earlier review of Isilon made the still valid point that optimized metadata structures for growing unstructured digital data also are key to accelerating delivery of digital content.<sup>1</sup>
3. There is a single point of management that is easy to use, yet powerful. The user interface is fully featured and built for control, not just management. Contrast this approach to that of many points of management often found in traditional SAN or NAS installation with multiple silos of storage.
4. Growth increments (both processing power and storage space) are reasonably sized and can be added without taking down the application. Capacity-on-demand pricing helps users manage operating expenses by allowing them to purchase only what is needed and only when it is needed.
5. The Isilon solution set is highly extensible and can be used alone or in combination with other EMC Big Data solutions. (See also point 7 below.)
6. This technology is forward-looking and consistent with overall trends in the IT industry. By acquiring Isilon, EMC brings the advantages of market power along with deeper pockets to fund future R&D investments. Isilon users can expect that their vendor will be able to respond to their product enhancement requests in a timely manner.
7. It is reassuring to current and future users that EMC offers more than just one solution to Big Data challenges. Isilon is just one part of the EMC solutions that leverage Big Data including *ATMOS* the online archival, backup, and restore tool for data in the cloud, the *Greenplum* product family of purpose-fit

<sup>1</sup> See [The Clipper Group Navigator](http://www.clipper.com/research/TCG2003069.pdf) dated December 17, 2003, entitled “*Isilon Accelerates Delivery of Digital Content*” and available at <http://www.clipper.com/research/TCG2003069.pdf>.

appliances for data analysis, and the widely publicized statement of EMC's support for *Apache Hadoop* open-source software for distributing, integrating, and supporting data-intensive distributed applications.<sup>2</sup>

Apache Hadoop has emerged as the preferred solution for Big Data analytics across unstructured data. EMC's support for Hadoop includes the availability of *EMC Greenplum HD Community Edition* and *EMC Greenplum HD enterprise edition*. This is a complete Hadoop solution including installation, training, and global support beyond simple packaging of the Apache Hadoop distribution.

Isilon can be used alone or in combination with other EMC-provided Big Data solutions.

8. The Isilon suite of products is fully functional and seems to have no obvious holes at this point.
9. Isilon customers span multiple industries thus proving the applicability of Isilon's solutions in virtually any domain that is experiencing huge data growth and/or a significant amount of unstructured data.
10. EMC has made a corporate commitment to being a leader in this market arena. It clearly is "not done", as evidenced by statements from Pat Gelsinger, President and COO, EMC Information Infrastructure Products. It has been widely reported in the commercial press that EMC is planning to spend \$3 billion on "Big Data" acquisitions over the near and medium term. Isilon and all of the other EMC Big Data solutions will be leveraged because EMC has made "Big Data" a corporate priority. Early adopters will appreciate that their data growth plans are matched by EMC's readiness to serve them.

### **Why This is Important**

At first glance, you may have had the impression that "Big Data" and Isilon's solutions to "Big Data" problems are not for you. After all, you likely are not exploring for oil and gas or doing biometric research or transmitting and interpreting MRIs. However, you may have already started down the "Big Data Path" and are just now realizing that fact. Consider whether your organization is performing any of the following activities or behaving in any of the following ways:

- Saving and accumulating a very large number of graphical images or photos.
- Performing repeated analytical procedures on, or modeling with, accumulated data.

- Never discarding transactional data, no matter how seemingly mundane it is.
- Capturing hits and history for your new social networking initiatives for possible later data mining.
- Using surveillance systems for site security and then saving the images.
- Saving web logs and email histories in anticipation of litigation defense strategies.
- Mixing traditional business data (text and numerical) with other forms of data (such as surveillance images) in order to discern insights not available when examining each data type alone.
- Saving call detail reporting for possible switch to different vendor or different technology or preparing it for data mining applications.
- Participating in any large-scale commerce activities, such as Internet sales, searches, or publishing and/or high -volume transaction processing.
- Using many RFID (Radio Frequency Identification) devices for inventory identification, shipping, or tracking.

### **Conclusion**

The (simple) truth is that almost every organization is capturing and saving a whole lot of data (some of it structured, but much unstructured) on the increasingly likely chance that they may need it someday. It is for this very reason that you should take a look into the future and see what kind of storage and management tools are available now and will be evolving in the near and medium term. One of the best examples for you to consider and investigate is Isilon from EMC. You really ought to put them on your radar screen – now – because you likely will need to solve your very own "Big Data Dilemma" far sooner than you may expect. Begin by taking inventory of what is now on your plate and how it would be managed in an Isilon or like environment. The effects will be surprising and enlightening and certainly will assist you in your preparation of longer-term plans. You even may decide that now is the time to invest in an EMC Isilon POC (Proof of Concept) system. Better get to it!



<sup>2</sup> See <http://www.emc.com/about/news/press/2011/20110509-03.htm>.

### ***About The Clipper Group, Inc.***

**The Clipper Group, Inc.**, is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at [www.clipper.com](http://www.clipper.com).***

### ***About the Author***

**Jim Baker is a Senior Contributing Analyst for The Clipper Group.** Mr. Baker specializes in storage hardware and software, the relationship between technology and its application in business, and emerging technological trends affecting users and vendors in the Information Technology community. He recently joined The Clipper Group after three decades in hardware and software product marketing and product management at companies as large as EMC and Honeywell Information Systems and as small as IPL Systems and Datapoint Corporation. He most recently served as Research Manager for Storage Software at IDC in Framingham, MA. Mr. Baker earned a Bachelor of Science in Marketing degree from the University of Illinois at Urbana-Champaign and an MBA, also from UIUC.

- **Reach Jim Baker via e-mail at [jim.baker@clipper.com](mailto:jim.baker@clipper.com) or at 781-235-0085 Ext. 212. (Please dial “212” when you hear the automated attendant.)**

### ***Regarding Trademarks and Service Marks***

**The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes,** and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***Disclosures***

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc. has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

### ***Regarding the Information in this Issue***

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc. cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.